

YANG TRISTA CAO

University of Maryland, College Park

ycao95@umd.edu ◊ (610)883-2633

<https://tristaycao.github.io/>

EDUCATION

University of Maryland, College Park

Ph.D Computer Science

Focus: Natural Language Processing

Advisor: Hal Daumé III

August 2018 - Present

Overall GPA: 3.81

Bryn Mawr College

Bachelor's of Arts degree in Computer Science

Concurrent major in Mathematics at Haverford College

Thesis Advisor: Lynne Butler

August 2014 - May 2018

Overall GPA: 3.72

RESEARCH INTEREST

My research lies in the intersection of Natural Language Processing and Human-Centered AI. I am interested in developing and evaluating natural language processing models to move towards trustworthy and fair NLP systems. Specifically, the research topics I am interested in are: fairness, trustworthy NLP, stereotypes, accessibility, visual question answering (VQA), multimodality, causality, explainable AI, trust, and human-centered AI.

PUBLICATIONS

What's Different between Visual Question Answering for Machine "Understanding" Versus for Accessibility?

Yang Trista Cao*, Kyle Seelman*, Kyungjun Lee*, Hal Daumé III

AAACL IJCNLP 2022 Best Theme Paper

Theory-Grounded Measurement of US Social Stereotypes in English Language Models

Yang Trista Cao*, Anna Sotnikova*, Hal Daumé III, Rachel Rudinger, Linda Zou

NAACL 2022 Long Paper

On the Intrinsic and Extrinsic Fairness Evaluation Metrics for Contextualized Language Representations

Yang Trista Cao*, Yada Pruksachatkun*, Kai-Wei Chang, Rahul Gupta, Varun Kumar, Jwala Dhamala, Aram Galstyan

ACL 2022 Short Paper

Analyzing Stereotypes in Generative Text Inference Tasks

Yang Trista Cao*, Anna Sotnikova*, Hal Daumé III, Rachel Rudinger

ACL-IJCNLP 2021 Findings Paper

Toward Gender-Inclusive Coreference Resolution:

An Analysis of Gender and Bias throughout the Machine Learning Lifecycle

Yang Trista Cao, Hal Daumé III

Computational Linguistics (preliminary version published at ACL 2019)

Controlling the Specificity of Clarification Question Generation

Yang Trista Cao, Sudha Rao, Hal Daumé III

Proceedings of the 2019 Workshop on Widening NLP

WORK EXPERIENCE

Amazon Research Intern, Cambridge, USA

May 2021 - August 2021

Mentor: Rahul Gupta and Kaiwei Chang

Project: On the Intrinsic and Extrinsic Fairness Evaluation Metrics for Contextualized Language Representations

AWARDS AND SCHOLARSHIPS

Best Theme Paper Award, AACL IJCNLP 2022.

Dean's Fellowship, University of Maryland, 2018-2020.

SERVICE

Co-organizer, the TrustNLP workshop, ACL 2022-2023.

Co-mentor, Welcoming Women to Computing Research, Tech + Research 2021.